

Discovering Knowledge using a Constraint-based Language

Patrice Boizumault, Bruno Crémilleux, Mehdi Khiari,
Samir Loudni, and Jean-Philippe Métivier

University of Caen Basse-Normandie – GREYC (CNRS UMR 6072)
Campus II Cote de Nacre, 14000 Caen - France
`{firstname.lastname}@unicaen.fr`

DPA 11201 – This work was presented into the Dagstuhl Seminar ”Constraint Programming meets Machine Learning and Data Mining” organized by Luc De Raedt, Heikki Mannila, Barry O’Sullivan, and Pascal Van Hentenryck - May 15-20, 2011.

Abstract. Discovering pattern sets or global patterns is an attractive issue from the pattern mining community in order to provide useful information. By combining local patterns satisfying a joint meaning, this approach produces patterns of higher level and thus more useful for the data analyst than the usual local patterns, while reducing the number of patterns. In parallel, recent works investigating relationships between data mining and constraint programming (CP) show that the CP paradigm is a nice framework to model and mine such patterns in a declarative and generic way. We present a constraint-based language which enables us to define queries addressing patterns sets and global patterns. The usefulness of such a declarative approach is highlighted by several examples coming from the clustering based on associations. This language has been implemented in the CP framework.

1 Introduction

Over the two last decades, local pattern discovery has become a rapidly growing field [16] and several paradigms are available for producing extensive collections of patterns such as the constraint-based pattern mining [17], condensed representations of patterns [3], interestingness measures [7] as well as integrating external resources and background knowledge [15]. Because of the exhaustive nature of the techniques, the pattern collections provide a fairly complete picture of the information content of the data. However, this approach suffers from limitations. First, the collections of patterns still remain too large for an individual and global analysis performed by the data analyst. Secondly, the so-called local patterns represent fragmented information and patterns expected by the data

analyst require to consider simultaneously several local patterns. In this work, we propose a declarative approach addressing the issue of discovering patterns combining several local patterns.

The data mining literature includes many methods to take into account the relationships between patterns and produce global patterns or pattern sets [4,8]. Recent approaches - constraint-based pattern set mining [4], pattern teams [14] and selecting patterns according to the added value of a new pattern given the currently selected patterns [2] - aim at reducing the redundancy by selecting patterns from the initial large set of local patterns on the basis of their usefulness in the context of the other selected patterns. Even if these approaches explicitly compare patterns, they are mainly based on the reduction of the redundancy or specific aims such as classification processes. Heuristic functions are often used and the lack of methods to mine complete and correct pattern sets or global patterns may be explained by the difficulty of the task. Mining local patterns under constraints requires the exploration of a large search space but mining global patterns under constraints is even harder because we have to take into account and compare the solutions satisfying each pattern involved in the constraints. The lack of generic approaches restrains the discovery of useful global patterns because the user has to develop a new method each time he wants to extract a new kind of global patterns. It explains why this issue deserves our attention.

In this paper, we propose a constraint-based language to discover patterns combining several local patterns. The data analyst expresses his/her queries thanks to constraints over terms built from constants, variables, operators, and function symbols. The key idea is to propose a generic and declarative approach to ask queries: the user models a problem by specifying a set of constraints and then a Constraint Programming (CP) system is responsible for solving it. This work is in the spirit of the cross-fertilization between data mining and CP which is a research field in emergence [10,11,12,13,18,19].

The constraint-based language offers the great advantage to provide an easy method to address different problems: it is enough to change the declarative specification in term of constraints. We illustrate the approach by several examples coming from the clustering based on associations: with simple query refinements, the data analyst is able to easily produce clusterings satisfying different properties. We think that the process greatly facilitates the building of global patterns and the discovery of knowledge. We do not detail in this paper the solving step, a preliminary implementation of the constraint-based language is given in [12].

This paper is organized as follows. Section 2 describes the constraint-based language and shows how queries and constraints can be defined using terms and built-in constraints. Starting from the clustering example, Section 3 depicts the process of successive refinements which enables us to easily address several kinds of clustering and then the discovery of global models.

Trans.	Items			
t_1	A	D	F	
t_2	A		E F	
t_3	A		E	G
t_4	A		E	G
t_5		B	E	G
t_6		B	E	G
t_7		C	E	G
t_8		C	E	G
t_9		C	E	H
t_{10}		C	E	H
t_{11}		C	F G H	

Table 1. Transactional dataset \mathcal{T} .

2 A Constraint-based Language

In this section, we describe the constraint-based language we propose. Terms are built using constants, variables, operators, and function symbols. Constraints are relations over terms that can be satisfied or not. First, we recall definitions. Then, we describe terms and present how the data analyst can define new function symbols using operators and built-in function symbols. Finally, we introduce constraints and show how queries and constraints can be defined using terms and built-in constraints.

2.1 Definitions and example

Let \mathcal{I} be a set of n distinct literals called items, an itemset (or *pattern*) is a non-null subset of \mathcal{I} . The language of itemsets corresponds to $\mathcal{L}_{\mathcal{I}} = 2^{\mathcal{I}} \setminus \emptyset$. A *transactional dataset* is a multi-set of m itemsets of $\mathcal{L}_{\mathcal{I}}$. Each itemset, usually called a *transaction* or object, is a database entry. For instance, Table 1 gives a transactional dataset \mathcal{T} where $m=11$ transactions t_1, \dots, t_{11} are described by $n=8$ items A, B, C, D, E, F, G, H .

Definition 1. (frequency) The frequency of a pattern is the number of transactions it covers. Let X_i be a pattern, $\mathbf{freq}(X_i) = |\{t \in \mathcal{T} \mid X_i \subseteq t\}|$.

So, $\mathbf{freq}(\{A, E\}) = 3$ and $\mathbf{freq}(\{C, F, G, H\}) = 1$. The *frequency* constraint focuses on patterns occurring in the dataset a number of times exceeding a given minimal threshold: $\mathbf{freq}(X_i) \geq \mathit{minfr}$. An other interesting measure to evaluate the relevance of patterns is the area [6].

Definition 2. (area) Let X_i be a pattern, $\mathbf{area}(X_i) = \mathbf{freq}(X_i) \times \mathbf{size}(X_i)$ where $\mathbf{size}(X_i)$ denotes the cardinality of X_i .

For transactional dataset \mathcal{T} (see Table 1), there are nine patterns satisfying the constraint $\mathbf{area}(X) \geq 6$: $\{A, E, G\}$, $\{B, E, G\}$, $\{C, E, G\}$, $\{C, E, H\}$, $\{E, G\}$, $\{C, E\}$, $\{C, H\}$, $\{E\}$, $\{G\}$.

2.2 Terms

Terms are built using:

1. **constants** are either numerical values (as threshold $minfr$), or items (as A) or patterns (as $\{A, B\}$) or transactions (as t_7).
2. **variables**, noted X_i , for $1 \leq i \leq k$, represent the unknown patterns.
3. **operators**:
 - set operators as $\cap, \cup, \setminus, \dots$
 - numerical operators as $+, -, \times, /, \dots$
4. **function symbols** involving one or several patterns: $\text{freq}/1, \text{size}/1, \text{cover}/1, \text{overlapItems}/2, \text{overlapTransactions}/2, \dots$

Terms are built using constants, variables, operators, and function symbols. Examples of terms:

- $\text{freq}(X_1) \times \text{size}(X_1)$
- $\text{freq}(X_1 \cup X_2) \times \text{size}(X_1 \cap X_2)$
- $\text{freq}(X_1) - \text{freq}(X_2)$

i) Built-in function symbols. Our constraint based language owns predefined (built-in) function symbols¹ like:

- $\text{cover}(X_i) = \{t \mid t \in \mathcal{T}, X_i \subseteq t\}$ is the set of transactions covered by X_i .
- $\text{freq}(X_i) = |\{t \mid t \in \mathcal{T}, X_i \subseteq t\}|$
- $\text{size}(X_i) = |\{j \mid j \in \mathcal{I}, j \in X_i\}|$
- $\text{overlapItems}(X_i, X_j) = |X_i \cap X_j|$ is the number of items shared by both X_i and X_j .
- $\text{overlapTransactions}(X_i, X_j) = |\text{cover}(X_i) \cap \text{cover}(X_j)|$ is the number of transactions covered by both X_i and X_j .

ii) User-defined function symbols. The data analyst can define new function symbols using constants, variables, operators and existing function symbols (built-in or previously defined ones). Examples:

- $\text{area}(X_i) = \text{freq}(X_i) \times \text{size}(X_i)$
- $\text{coverage}(X_i, X_j) = \text{freq}(X_i \cup X_j) \times \text{size}(X_i \cap X_j)$
- Let $D_1, D_2 \subset \mathcal{T}$ be 2 sets of transactions and $\text{freq}(X_i, D_j)$ the frequency of pattern X_i into D_j , then:

$$\text{growth-rate}(X_i) = \frac{|D_2| \times \text{freq}(X_i, D_1)}{|D_1| \times \text{freq}(X_i, D_2)}$$

¹ Only function symbols used in Section 3 are introduced in this paper.

2.3 Constraints and Queries

Constraints are relations over terms. They can be either *built-in* or *user-defined*. There are three kinds of constraints:

1. **numerical** ones like: $<, \leq, =, \neq, \geq, >, \dots$

Examples:

- $\text{freq}(X_1) \leq 10$
- $\text{size}(X_2) = 2 \times \text{size}(X_3)$
- $\text{area}(X_1) < \text{size}(X_2) \times \text{size}(X_3)$

2. **set** ones like: $=, \neq, \in, \notin, \subset, \subseteq, \dots$

Examples:

- $i_3 \in X_1$
- $X_1 \cup X_2 \subset X_3$
- $X_1 = X_2 \cap X_4$

3. **dedicated** ones like:

- $\text{closed}(X_i)$ is satisfied iff X_i is a closed² pattern.
- $\text{coverTransactions}([X_1, \dots, X_k])$ is satisfied iff each transaction is covered by at least one pattern (i.e. $\bigcup_{1 \leq i \leq k} \text{cover}(X_i) = \mathcal{T}$),
- $\text{coverItems}([X_1, \dots, X_k])$ is satisfied iff every item belongs to at least one pattern (i.e. $\bigcup_{1 \leq i \leq k} X_i = \mathcal{I}$).
- $\text{canonical}([X_1, \dots, X_k])$ is satisfied iff for all i s.t. $1 \leq i < k$, pattern X_i is less than pattern X_{i+1} with respect to the lexicographic order.

Queries and constraints are formulae built using constraints and logical connectors: \wedge (conjunction) and \vee (disjunction).

In the following, we take the exception rules as example³. An *exception rule*⁴ is a pattern combining a strong rule and a deviational pattern to the strong rule:

$$e(X_1 \rightarrow \neg I) \equiv \begin{cases} \text{true} & \text{if } \exists X_2 \in \mathcal{L}_{\mathcal{I}} \text{ such that } X_2 \subset X_1, \text{ one have} \\ & (X_1 \setminus X_2 \rightarrow I) \wedge (X_1 \rightarrow \neg I) \\ \text{false} & \text{otherwise} \end{cases}$$

adding X_2 to $X_1 \setminus X_2$ provides the exception rule $X_1 \rightarrow \neg I$

- $X_1 \setminus X_2 \rightarrow I$ must be a frequent rule having a high confidence value:
- $X_1 \rightarrow \neg I$ must be a rare rule having a high confidence value:

to sum up:

$$\text{exception}(X_1, X_2) \equiv \begin{cases} \text{freq}((X_1 \setminus X_2) \sqcup I) \geq \text{minfr} \wedge \\ (\text{freq}(X_1 \setminus X_2) - \text{freq}((X_1 \setminus X_2) \sqcup I)) \leq \delta_1 \wedge \\ \text{freq}(X_1 \sqcup \neg I) \leq \text{maxfr} \wedge \\ (\text{freq}(X_1) - \text{freq}(X_1 \sqcup \neg I)) \leq \delta_2 \end{cases}$$

² Let Tr_i be the set of transactions covered by pattern X_i . X_i is closed iff X_i is the largest (\subset) pattern covering Tr_i .

³ For more examples, see the modelling of the clustering problem (Section 3).

⁴ The definition of exception rules initially presented in [20] also includes a reference rule $X_2 \not\rightarrow \neg I$.

3 From Modelling to Solving

The major strength of our approach is to provide a simple and efficient way to refine a query. In practice, the data analyst begins with submitting a first query Q_0 . Then, he will successively refine this query (deriving Q_{i+1} from Q_i) until he considers that relevant information has been extracted.

Clustering models aim at partitioning data into groups (clusters) so that transactions occurring in the same cluster are similar but different from those appearing in other clusters. We selected the clustering problem to illustrate our approach for two main reasons. First, clustering is an important and popular unsupervised learning method [1,5,9]. Then, by nature, clustering proceeds by iteratively refining queries until a satisfactory solution is found. The clustering model, used here, starts from closed patterns because a closed pattern is a pattern gathering the maximum amount of similarity between a set of transactions.

3.1 Modelling a clustering query

The usual clustering problem can be defined as follows:

“to find a set of k closed patterns X_1, X_2, \dots, X_k covering all transactions without any overlap on these transactions”.

First, **closed**(X_i) constraints (see Section 2.3) are used to enforce each unknown pattern X_i to be closed,

Then, it is easy to constrain the set of patterns to cover all the transactional dataset using the **coverTransactions**[X_1, X_2, \dots, X_k] constraint (see Section 2.3).

Finally, to avoid any overlap over the transactions, for each couple of patterns (X_i, X_j), $i < j$, a constraint **overlapTransactions**(X_i, X_j) = 0 is added. This constraint states that there is no transaction covered by both X_i and X_j .

The following query (Q_0) models the initial clustering problem:

$$\begin{cases} \bigwedge_{1 \leq i \leq k} \text{closed}(X_i) \wedge \\ \text{coverTransaction}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapTransactions}(X_i, X_j) = 0 \end{cases}$$

On our running example, when looking for a clustering with $k = 3$ patterns, we obtain 30 solutions (See Table 2).

3.2 Refining queries

By only refining queries addressing a clustering, the data analyst can easily produce clusterings satisfying different properties. In this section, we illustrate this approach by successive refinements. Starting from initial query Q_0 , symmetrical solutions are first removed leading to query Q_1 . Then, clusterings with non-frequent patterns and clusterings with small size patterns are removed (leading to queries Q_2 and Q_3). More generally, this process greatly facilitates the building of global patterns and the discovery of knowledge.

Sol.	X_1	X_2	X_3
s_1	{C, F, G, H}	{E}	{A, D, F}
s_2	{C, F, G, H}	{A, D, F}	{E}
s_3	{A, D, F}	{C, F, G, H}	{E}
s_4	{A, D, F}	{E}	{C, F, G, H}
s_5	{E}	{C, F, G, H}	{A, D, F}
s_6	{E}	{A, D, F}	{C, F, G, H}
s_7	{A, F}	{C, H}	{E, G}
\vdots	\vdots	\vdots	\vdots
s_{13}	{C, E, H}	{E, G}	{F}
\vdots	\vdots	\vdots	\vdots
s_{19}	{A, F}	{C, E, H}	{G}
\vdots	\vdots	\vdots	\vdots
s_{25}	{A}	{B, E, G}	{C}
\vdots	\vdots	\vdots	\vdots
s_{30}	{C}	{B, E, G}	{A}

Table 2. Set of all solutions (including symmetrical ones).

i) Removing symmetrical solutions. Two solutions s_i and s_j are said to be symmetrical iff there exists a permutation σ , such that $s_j = \sigma(s_i)$. A clustering problem owns intrinsically a lot of symmetrical solutions: let $s = (p_1, p_2, \dots, p_k)$ be a solution containing k patterns p_i . Any permutation σ of these k patterns $\sigma(s) = (p_{\sigma(1)}, p_{\sigma(2)}, \dots, p_{\sigma(k)})$ is also a solution. So, for any solution, there exist $(k! - 1)$ symmetrical solutions. For example, solutions from s_1 to s_6 are symmetrical (See Table 2) and constitute the same clustering.

Constraint **canonical**($[X_1, \dots, X_k]$) is used to avoid symmetrical solutions. This constraint states that, for all i s.t. $1 \leq i < k$, pattern X_i is less than pattern X_{i+1} with respect to the lexicographic order.

From query Q_0 , we obtain query Q_1 :

$$\begin{cases} \bigwedge_{1 \leq i \leq k} \text{closed}(X_i) \wedge \\ \text{coverTransaction}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapTransactions}(X_i, X_j) = 0 \wedge \\ \text{canonical}([X_1, \dots, X_k]) \end{cases}$$

Following our running example, query Q_1 leads to only 5 solutions since $5 \times 3! = 30$ (See Table 3).

The constraint **canonical**($[X_1, \dots, X_k]$) plays an important role. First, as the number of solutions $(k!)$ grows very rapidly with the number k of clusters, it

Sol.	X_1	X_2	X_3
s_1	{C, F, G, H}	{E}	{A, D, F}
s_7	{A, F}	{C, H}	{E, G}
s_{13}	{C, E, H}	{E, G}	{F}
s_{19}	{A, F}	{C, E, H}	{G}
s_{25}	{A}	{B, E, G}	{C}

Table 3. Set of different clusterings.

quickly becomes very large. So, it is essential and indispensable to break the symmetries to avoid having a huge number of redundant solutions. Moreover, this constraint will perform an efficient filtering by drastically reducing the size of the search space.

ii) Removing solutions with non-frequent patterns. A clustering containing at least one pattern having a low frequency is not considered to be relevant. To remove such solutions, we only need to add new constraints to the current query Q_1 . Such a constraint requires that each cluster must have a frequency greater than a threshold (here 10% of $m = 11$).

$$\forall 1 \leq i \leq k, \text{freq}(X_i) \geq 2$$

From query Q_1 , we obtain query Q_2 :

$$\left\{ \begin{array}{l} \wedge_{1 \leq i \leq k} \text{closed}(X_i) \wedge \\ \text{coverTransaction}([X_1, \dots, X_k]) \wedge \\ \wedge_{1 \leq i < j \leq k} \text{overlapTransactions}(X_i, X_j) = 0 \wedge \\ \text{canonical}([X_1, \dots, X_k]) \wedge \\ \wedge_{1 \leq i \leq k} \text{freq}(X_i) \geq 2 \end{array} \right.$$

Pattern $\{C, F, G, H\}$ of solution s_1 (see Table 2) has a frequency of 1 which is less than the threshold. So for Q_2 , solution s_1 is not valid. For query Q_2 , there remain 4 solutions: s_7 , s_{13} , s_{19} , and s_{25} (See Table 3).

iii) Removing solutions with small size patterns. A clustering containing at least one pattern of size 1 is not considered to be relevant⁵. To remove such clusterings, we only need to add new constraints to the current query Q_2 . Such a constraint requires that each cluster must have a size greater than 1. This can be achieved by stating, for each cluster, a constraint to restrict its size.

$$\forall 1 \leq i \leq k, \text{size}(X_i) \geq 2$$

⁵ Usually, clusterings using these unitary clusters reflect the discretisation of some attributes.

From query Q_2 , we obtain query Q_3 :

$$\begin{cases} \bigwedge_{1 \leq i \leq k} \text{closed}(X_i) \wedge \\ \text{coverTransaction}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapTransactions}(X_i, X_j) = 0 \wedge \\ \text{canonical}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i \leq k} \text{freq}(X_i) \geq 2 \wedge \\ \bigwedge_{1 \leq i \leq k} \text{size}(X_i) \geq 2 \end{cases}$$

Query Q_3 has only 1 solution: s_7 (see Table 3). For this solution, we have $X_1 = \{A, F\}$, $X_2 = \{C, H\}$ and $X_3 = \{E, G\}$.

3.3 Solving other Clustering Problems

In the same way, it is easy to express other clustering problems such as co-clustering, soft clustering and soft co-clustering.

i) **The soft clustering problem** is a relaxed version of the clustering problem where small overlaps (less than δ_T) on transactions are authorized. This problem is modelised by query Q_4 (soft version of Q_0):

$$\begin{cases} \bigwedge_{1 \leq i \leq k} \text{closed}(X_i) \wedge \\ \text{coverTransaction}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapTransactions}(X_i, X_j) \leq \delta_T \wedge \\ \text{canonical}([X_1, \dots, X_k]) \end{cases}$$

Consider query Q_4 with $k=3$ and a maximal overlap for transactions $\delta_T=1$. There are 13 solutions (see Table 4). If symmetries are not broken using the constraint $\text{canonical}([X_1, \dots, X_k])$, then there are 78 ($3! \times 13$) solutions.

For solution s'_1 , patterns X_1 and X_3 cover transaction t_{11} (see Table 1). Moreover, patterns X_2 and X_3 cover transaction t_2 (see Table 1). After having removed solutions with non-frequent patterns, there remain 8 solutions: from s'_6 to s'_{13} . After having removed solutions with small size patterns, it remains only 1 solution: s'_9 (which is the solution s_7 of the initial clustering problem, see Section 3.1).

ii) **The co-clustering problem** consists in finding k clusters covering both the set of transactions and the set of items, without any overlap on transactions or on items. This problem is modelised by query Q_5 :

$$\begin{cases} \bigwedge_{1 \leq i \leq k} \text{closed}(X_i) \wedge \\ \text{coverTransaction}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapTransactions}(X_i, X_j) = 0 \wedge \\ \text{coverItems}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapItems}(X_i, X_j) = 0 \wedge \\ \text{canonical}([X_1, \dots, X_k]) \end{cases}$$

Sol.	X_1	X_2	X_3
s'_1	{C, F, G, H}	{E}	{F}
s'_2	{A, D, F}	{C, F, G, H}	{E}
s'_3	{A, F}	{C, F, G, H}	{E}
s'_4	{A, E, F}	{E}	{F}
s'_5	{A, D, F}	{E}	{F}
s'_6	{A}	{B, E, G}	{C}
s'_7	{A, F}	{C, E, H}	{G}
s'_8	{A, F}	{C, H}	{G}
s'_9	{A, F}	{C, H}	{E, G}
s'_{10}	{C, E, H}	{E, G}	{F}
s'_{11}	{C, H}	{E, G}	{F}
s'_{12}	{C, H}	{F}	{G}
s'_{13}	{C, E, H}	{F}	{G}

Table 4. Set of different clusterings for query Q_4 (soft clustering).

iii) **The soft co-clustering problem** is a relaxed version of the co-clustering problem, allowing small overlaps on transactions (less than δ_T) and on items (less than δ_I). This problem is modelised by query Q_6 (soft version of Q_4 and Q_5):

$$\left\{ \begin{array}{l} \bigwedge_{1 \leq i \leq k} \text{closed}(X_i) \wedge \\ \text{coverTransaction}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapTransactions}(X_i, X_j) \leq \delta_T \wedge \\ \text{coverItems}([X_1, \dots, X_k]) \wedge \\ \bigwedge_{1 \leq i < j \leq k} \text{overlapItems}(X_i, X_j) \leq \delta_I \wedge \\ \text{canonical}([X_1, \dots, X_k]) \end{array} \right.$$

4 Conclusions and Future Works

We have proposed a constraint-based language allowing to easily express different mining tasks in a declarative way. Thanks to the declarative process, extending or changing the specification to refine the results and get more relevant patterns or address new global patterns is very simple. Moreover, all constraints can be combined together and new constraints can be added.

The effectiveness and the flexibility of our approach is shown on several examples coming from clustering based on associations: thanks to query refinements, the data analyst is able to produce clusterings satisfying different constraints, thus generating more meaningful clusters and avoiding outlier ones.

As future work, we intend to enrich our constraint-based language with further constraints to capture and model a wide range of data mining tasks. The scalability of the approach to larger values of k and larger datasets can also be investigated. Another promising direction is to integrate optimisation criteria in our framework.

References

1. Pavel Berkhin. Survey of clustering data mining techniques. Technical report, Accrue Software, San Jose, CA, USA, 2002.
2. B. Bringmann and A. Zimmermann. The chosen few: On identifying valuable patterns. In *proceedings of the 12th IEEE International Conference on Data Mining (ICDM-07)*, pages 63–72, Omaha, NE, 2007.
3. T. Calders, C. Rigotti, and J-F. Boulicaut. A survey on condensed representations for frequent sets. In J-F. Boulicaut, L. De Raedt, and H. Mannila, editors, *Constraint-Based Mining and Inductive Databases*, volume 3848 of *LNCS*, pages 64–80. Springer, 2005.
4. L. De Raedt and A. Zimmermann. Constraint-based pattern set mining. In *Proceedings of the Seventh SIAM International Conference on Data Mining*, Minneapolis, Minnesota, USA, April 2007. SIAM.
5. Douglas H. Fisher. Knowledge acquisition via incremental conceptual clustering. *Machine Learning*, 2(2):139–172, 1987.
6. Floris Geerts, Bart Goethals, and Taneli Mielikäinen. Tiling databases. In *Discovery Science*, volume 3245 of *LNCS*, pages 278–289. Springer, 2004.
7. L. Geng and H. J. Hamilton. Interestingness measures for data mining: A survey. *ACM Comput. Surv.*, 38(3), 2006.
8. A. Giacometti, E. Khanjari Miyaneh, P. Marcel, and A. Soulet. A framework for pattern-based global models. In *10th Int. Conf. on Intelligent Data Engineering and Automated Learning (IDEAL'09)*, volume 5788 of *LNCS*, pages 433–440, Burgos, Spain, 2009.
9. David Gibson, Jon M. Kleinberg, and Prabhakar Raghavan. Clustering categorical data: An approach based on dynamical systems. In Ashish Gupta, Oded Shmueli, and Jennifer Widom, editors, *VLDB*, pages 311–322. Morgan Kaufmann, 1998.
10. Tias Guns, Siegfried Nijssen, and Luc De Raedt. Itemset mining: A constraint programming perspective. *Artif. Intell.*, 175(12-13):1951–1983, 2011.
11. M. Khiari, P. Boizumault, and B. Crémilleux. Local constraint-based mining and set constraint programming for pattern discovery. In *From Local Patterns to Global Models (LeGo-09), ECML/PKDD-09 Workshop*, pages 61–76, Bled, Slovenia, 2009.
12. M. Khiari, P. Boizumault, and B. Crémilleux. Constraint programming for mining n-ary patterns. In *16th International Conference on Principles and Practice of Constraint Programming (CP'10)*, volume 6308 of *LNCS*, pages 552–567, St Andrews, Scotland, 2010. Springer-Verlag.
13. M. Khiari, P. Boizumault, and B. Crémilleux. Combining CSP and constraint-based mining for pattern discovery. In *Advances in Knowledge Discovery and Management (Post-EGC Selected Papers)*. Springer-Verlag, 2011. 20 pages.
14. A. Knobbe and E. Ho. Pattern teams. In J. Fürnkranz, T. Scheffer, and M. Spiliopoulou, editors, *proceedings of the 10th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'06)*, volume 4213 of *LNAI*, pages 577–584, Berlin, Germany, September 2006. Springer-Verlag.
15. Nada Lavrac, Filip Zelezny, and Saso Dzeroski. Local patterns: Theory and practice of constraint-based relational subgroup discovery. In *Local Pattern Detection*, volume 3539 of *LNCS*, pages 71–88. Springer, 2005.
16. K. Morik, J-F. Boulicaut, and A. Siebes (eds.), editors. *Local Pattern Detection*, volume 3539 of *LNAI*. Springer-Verlag, 2005.
17. R. T. Ng, V. S. Lakshmanan, J. Han, and A. Pang. Exploratory mining and pruning optimizations of constrained associations rules. In *proceedings of ACM SIGMOD'98*, pages 13–24. ACM Press, 1998.

18. Luc De Raedt, Tias Guns, and Siegfried Nijssen. Constraint programming for itemset mining. In Ying Li, Bing Liu, and Sunita Sarawagi, editors, *KDD*, pages 204–212. ACM, 2008.
19. Luc De Raedt, Tias Guns, and Siegfried Nijssen. Constraint programming for data mining and machine learning. In *Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, pages 1671–1675, 2010.
20. E. Suzuki. Undirected discovery of interesting exception rules. *International Journal of Pattern Recognition and Artificial Intelligence*, 16(8):1065–1086, 2002.